

604949  
604949

20-0

|            |           |    |   |    |
|------------|-----------|----|---|----|
| COPY       | 1         | OF | 1 | 52 |
| ...RD COPY | \$ . 1.00 |    |   |    |
| MICROFICHE | \$ . 0.50 |    |   |    |




THE BACKGROUND AND IMPLICATIONS  
OF THE SYSTEMS RESEARCH LABORATORY STUDIES

Robert L. Chapman and John L. Kennedy

21 September 1955

P-740

Approved for OTS release



The RAND Corporation  
1700 MAIN ST. • SANTA MONICA • CALIFORNIA

### SUMMARY

The effective operation of complex man-machine systems <sup>problem</sup> is a major Air Force problem, and one that The RAND Corporation's Systems Research Laboratory has been working on for the last four years. Solutions to this problem require an understanding of group rather than individual behavior. In its laboratory studies of complete organizations the Systems Research Laboratory has used a full-scale model of an air-defense direction center. It has studied this model in great detail in much the same way as scale model aircraft are studied in wind tunnels—by manipulating an environment to apply stress so that weaknesses that show up in the model can be used as the basis for making improvements. But organizations change under stress because they learn, and <sup>was the independent of</sup> one major result of this study has been techniques for using group learning in greatly improving system performance.)

The characteristics of the air-defense direction center used as the subject of a study of organizational development simplified the problem of research and made it possible to see more of the development process. In dealing with group learning this organization was considered a unit rather than a group of individuals. Group learning can be described by quantitative task and response models. These models, together with quantitative descriptions of stress on the organization and of conditions that affect learning, can be used to predict system performance.

In addition to a theoretical approach that can be used by other scientists, the results of this study have many implications for system design, especially in personnel selection and training and in human engineering.

Paper presented at the Air Force—National Research  
Council Scientific Symposium, Washington, D. C.,  
November 14-15, 1955.

No Air Force mission can be accomplished without men.

The accomplishment of many Air Force missions depends on the ability of two or more individuals to act in cooperation as a group. . . Failure of weapon systems may be due to failure in effective crew functioning even when individual group members possess the requisite skills and abilities for their individual jobs.

. . . Increasing the effectiveness with which the Air Force can accomplish its peace and wartime missions by promoting better group functioning among individuals working together. . . [can be obtained] through developing means of evaluating the group's performance, of identifying the organization, [of finding] policies and procedures of the group which most directly contribute to superior performance, and of testing methods for assembly and training of groups which most likely will provide effective units.

This excellent statement of the problems of making organizations effective is from a recent Air Force Technical Document. It not only describes one of the Air Force's most critical problems, but in pointing out where solutions might be found it also defines the purpose of The RAND Corporation's Systems Research Laboratory. What this Laboratory is trying to do is precisely this development of "means of evaluating group performance," this "identifying the organization" and finding the best policies and procedures and the best training methods.

Our present technological society has become almost completely dependent on large, complex man-machine systems. But these systems have become so big and so complex that they are almost beyond the comprehension of the men who operate and direct them. As the systems have grown in size and complexity, the tasks of the men who run them have too. Bigger and more complicated machinery won't necessarily give us better results. It's necessary to understand the behavior of the men who operate these systems, and since

systems are run by teams and not by individuals, understanding the critical human elements of these systems means going beyond individual psychology into the terra incognita of organizational behavior.

The continental air-defense system is a good example of what's been happening. Just a few years ago an aircraft-warning net that surrounded the entire country would have been little more than a wild dream—a communication network of thousands of men and machines that linked together fighter bases, radar sites, interceptors, civilian-defense groups, the Civil Aeronautics Authority, and other organizations that had to work together on split-second schedules was patently out of the question. But this tremendous (and tremendously complicated) system is in operation today, and a bigger and more complex one will be operating tomorrow.

In using systems like this one there are new problems of selecting, training, and utilizing men--problems in which questions of the proper use of individual specialists are only a beginning. Today's problems have become those of group coordination and integration, of team performance and team learning, the problems that are the critical ones in understanding and operating man-machine systems.

Understanding how groups of men work together effectively is obviously of very general interest, in the industrial and commercial world as well as to the armed services, and a problem that requires a variety of research efforts. Operations analysts, industrial engineers, sociologists, political scientists, psychologists, planners, and executives—to mention only a few—have been working on it.

The RAND Corporation's Systems Research Laboratory has been studying these complex systems in a laboratory with experimental methods. The

Laboratory has been studying complete systems, not just parts of systems, and--probably most important--it has studied them as man-machine systems, deliberately considering men as integral parts of the system.

This paper will present a general picture of how this was done, what was learned by doing it, and what the important implications of these studies might be and where they will fit in the general scheme of research activity.

### THE SYSTEM STUDIED

It became obvious early in these studies that simulation techniques and techniques for controlling large-scale experiments would have to be pushed beyond their current state of development and that it would be necessary to select a kind of organization that lent itself to being studied with the new techniques. The system selected was suited to these laboratory techniques; it was also of critical importance to the Air Force and one that had enough in common with other systems to give the results generality.

#### The Air-Defense Direction Center

This system was the air-defense direction center, an organization that defends a portion of the United States against enemy air attack. In many ways a direction center is a complete system; it has all the information available about the air traffic in its area and controls weapons for stopping enemy air attacks. What was simulated in the laboratory was a close approximation to a real direction center--a full-scale model manned by a standard crew of 30 to 40 men. Four air-defense experiments were conducted. Each ran for about 200 hours--the equivalent of about six weeks of normal life in a real direction center.

Figures 1 and 2 (Crew Members in Laboratory)

A direction center is a rather complex organization with quite a complex job to do. The laboratory crews had to defend an area of roughly 100,000 square miles. During each experiment there were about 10,000 flights over this area. The air traffic, which increased more than threefold during the experiment, included a wide variety of flights—from commercial air liners on transoceanic flights to cub aircraft hedgehopping from airport to airport. Hostile attacks on targets in the area ranged from single bombers trying to camouflage themselves in the flight-plan traffic to mass raids of as many as 25 hostiles. Symbols containing information about these flights came into the system at an average of 300 a minute—a rate of information input that added up to something like two million symbols during an experiment.

#### Choosing a System to Study

Two conceptual issues, resolved before the experiments started, were crucial ones for making these studies possible. Both of these issues, which were concerned with the kind of organization to be studied, delimited a complicated problem.

First of all, the air-defense direction center is an organization in which task accomplishment has a well-accepted social value and one whose successes and failures are fairly easy to evaluate at almost any time during its operation. An experimenter can have confidence in an air-defense crew's motivation to defend the country against air attack and in recognition of success and failure. And this motivation is complicated very little by previous personal histories. The complex of values, attitudes, and beliefs that influence this organization's development are derived mostly from the

crew's experience with air defense. Because the groups studied were newly assembled, a good part of this happened right in the laboratory.

A second advantage of studying a direction center is that more of the group's activity can be observed than in many other organizations. Much of the crew's behavior in dealing with its task is verbal response to known stimuli--either to other verbal behavior or to task information coming into the system. There is little of importance that can't be seen or heard by the experimenter.

Since these experiments involved groups of nearly 40 men, choosing a system that had these characteristics simplified the problem tremendously. The motivation of the men under study and the means of measuring system effectiveness were both relatively uncomplicated. Most of the group's relevant behavior--and the way this behavior changed--was exposed to view.

#### THE METHODS USED

Although a description of the system studied gives some idea of the size and scope of the experiments, the ideas behind these experiments can be put into a larger context and one that is probably more meaningful. Since the effects of equipment modifications were not the object of study, the physical resources were kept constant during each experiment and the task was varied. Any improvements in performance depended entirely on each crew's skill in using the resources it already had.

The Systems Research Laboratory's facilities are used to study human organizations in much the same way a wind tunnel is used in developing new aircraft. In both methods the experimenters manipulate an environment to apply stress to highly detailed models so that the performance of the prototype can be predicted and changes made to improve it. A wind tunnel



uses a detailed scale model of the aircraft whose flight characteristics are being studied; the Systems Research Laboratory used a model organization of 30 to 40 men that was practically full scale. By exposing the models to critical environmental conditions over and over again in different combinations, both facilities can be used to expose weak points in the design of the prototype.

Both wind tunnels and this way of studying organizations rely heavily on elaborate measuring devices. And both of them accumulate enormous amounts of data--so much, in fact, that a corps of specialized professional, technical, and clerical workers is needed to handle it.

Research facilities such as the Systems Research Laboratory, again like wind tunnels, are big and expensive, but they may well become as indispensable in designing and improving systems as wind tunnels are in designing aircraft.

But with all these similarities, there is one main difference between wind-tunnel studies of aircraft and large-scale laboratory research with human organizations. In experiments with organizations, the laboratory model changes under stress. It learns. Learning is an invaluable characteristic. It is also a complicating one. Because organizations learn, a formula for predicting their performance, unlike a formula for predicting the behavior of aircraft, has to take into account the way the organization changes under stress.

Although aspiring to study complete man-machine systems is obviously fine in principle, worthwhile results depend on how effectively aspirations are translated into experimental form. An important aspect of this

translation is gaining "observational access to the phenomena." "Observational access" is more than being able to get meaningful data -- it's primarily a problem of getting worthwhile phenomena to occur at all. If an organization is to be observed under a variety of conditions, it's essential that the men who are being studied function as an organization and not just as a group of individuals and that they are stimulated to develop as an organization--to learn as a group. This failed to happen in the first experiment--the organization learned so much faster than it had been expected to that long before the experiment was over the task that had been so carefully prepared became so easy that the group's performance was no longer worth observing.

#### RESULTS AND THEORY

The outstanding empirical result of these experiments was the degree to which an air-defense crew can learn to use its resources more effectively. That a group of human beings can learn is by no means a momentous conclusion--after all, it seems rather obvious to say that the performance of a system can be improved if it has resources of one kind or another that it hasn't used before. What was startling in these experiments was the extent to which performance could be improved by exploiting these unused resources. Although the task load was increased gradually so that it was more than three times as great at the end of the experiment as it had been in the beginning, each of the four crews kept up a highly effective defense of the area against enemy air attack.

Because an organization whose achievement is readily measured was chosen for study, the evidence for saying that organizational development

did take place is readily found. Although traffic was continually increased during each experiment to the point where, in the last part of each experiment, it was heavier than the normal air traffic in any part of the United States, each crew's defense against hostile attacks of all kinds continued at a more effective level than we had any reason to expect. (Incidentally there were so many similarities in performance and development among the four crews that crew learning can be considered in the singular, since what happened in any one of the crews was fairly typical of all of them.)

#### Task and Response Models

But the scientific significance of the Systems Research Laboratory's work is the way these experiments exposed the process of organizational development.

Just what does an air-defense crew do to maintain effective performance in dealing with a task that keeps getting harder and harder? A rather obvious answer is that it spends its efforts more efficiently. With each increase in the number of tracks the crew had to deal with saturation seemed imminent because the crew found it more and more difficult to continue handling each track with its current procedures. But each time that saturation seemed imminent, some way of simplifying the job was found.

One way to measure the effort a crew expends is by the number of items of information, such as position reports, it uses to handle the task. There was only a slight increase in the rate of information flow during an experiment. As a matter of fact, during the last hour of the experiment, when the load was more than three times as heavy, the crew used just about the same amount of information it did during the first hour.

It maintained this unexpectedly high degree of success in defending the area by concentrating on traffic that was potentially hostile, spending smaller and smaller amounts on the rest of the tracks. If the crew had spent its efforts at the same rate during the last hour as it did during the first, it would have used nearly 1,300 items of information. Actually, it used only 640—just about half of what would have been necessary if it hadn't changed its ways of handling the task. This is one example of more effective use of the same amount of effort—an illustration of how the crew assigns the kind and amount of effort to task events it considers important. This rough measure of effort expended is the "response model" (Fig. 3, page 9a).

But since there are so many task events, the crew must have some way of deciding which ones are important. It does this by making distinctions between tracks that it has to deal with to accomplish the task and those that it doesn't have to deal with at all. These progressively finer distinctions about which classes of tracks need to be handled make up the "task model" (Fig. 4, page 9a). Although the number of tracks in the task increased steadily, there was only a slight increase in the number of tracks the crew dealt with. Since it continued to defend the area successfully, even though it dealt with only part of the tracks (about 40 per cent of them in the last hour), these distinctions were obviously effective ones. The important discriminations were between threatening flights (traffic coming from certain directions) and nonthreatening ones (traffic going in other directions).

These models enable the organization to spend its effort more effectively by determining what efforts will be given priority. By making appropriate




Fig. 3. Although there is only a slight increase<sup>855 N</sup> in the amount of effort a crew spends during an experiment, more and more of it is spent on critical tracks.




Fig. 4. The number of tracks the crew handles does not increase as fast as the number of tracks in the stimulus (left). This occurs because it handles a smaller and smaller proportion of noncritical tracks (right).<sup>855 M</sup>

changes in the models the organization can adapt to changing task circumstances.

### A Theory of Organizational Learning

These empirical results seem to indicate that an organization will look for new patterns of behavior when it needs them—when it is under stress. Stress in an organization seems to arise from failure to perform effectively or—for an equally important reason—because it has to work too hard to avoid failure. The first is "failure" stress; the second, "discomfort" stress.

This effect of stress on organizations suggests an analogy between group learning and the familiar description of individual learning: stress, new and appropriate response, reinforcement. Without stress, organizations don't learn. Without reinforcement, they don't learn rapidly.

The results of these experiments indicate that group learning is an essential factor in any equation for predicting organizational effectiveness. From the analogy to individual learning, the main outlines of this theory seem clear—it must include the source of stress (the discomfort and failure that act as pressure to learn), and ways of reducing stress (the priority schemes of the task and response models). Such a formulation should help to predict how fast and how far a system can adapt, to identify what is difficult in the task, and to define the conditions that help an organization use its resources most effectively.

Perhaps the most important result of these experiments is that such concepts as stress, rate of learning, and so on can be described quantitatively. There are, to be sure, some practical difficulties—the Laboratory

now has over 12,000 hours of recordings and some 60 file drawers of supporting information from the air-defense studies. Thus far, some 100,000 IBM cards have been coded for each experiment, with perhaps an equally large amount of information, not yet successfully coded, left over.

These coded data are being used to represent measurable failure and discomfort stress, and the relationship between stress and changes in the task and response models is being explored. This has brought up questions of the place of energy expenditure in group learning and of the sequence of successive steps in adaptation. Adaptation seems to involve a complicated feedback process. When the task becomes more difficult, the crew absorbs some discomfort—making only those changes it can make readily—in the task model. But this expedient may well add failure to discomfort. Making further changes that are necessary—in the response model—requires a greater degree of coordination. These changes require additional skill, and the time needed to acquire the needed skill may be another source of failure. As the crew adapts to successive failure and discomfort in this way, the task and response models gradually stabilize, much as an oscillating function damps.

But adaptation is affected by many details not yet fully understood—external conditions in the environment and internal conditions in the organization itself that help or hinder learning. An example of these conditions is the "grease pencils are no damn good" symptom of stress. An organization's first reaction to stress seems to be to blame external conditions, faults of the equipment, and so on. At one point or another during the experiments each crew blamed ineffective operation on the grease

pencils it used for marking plots on the big movement board--the pencils were too hard or too soft, they broke too easily, they weren't the right color. But complaints like these disappear when the crew begins to find ways of doing things that lead to better performance.

The analogy between group learning and individual learning suggests the substitution of the organization for the individual as the organism in the classical learning model when organizational adaptation is being considered. In these air-defense experiments the organization has been treated as a unit rather than a collection of individuals, not only in managing the experimental conditions but also in analyzing the data and in building a theoretical framework on the basis of the results. With this kind of formulation the characteristics of individuals--their personality and skill--appear only as qualities of the organization. Such a formulation of group learning seems consistent with much of the data and has some rather definite implications.

#### SOME IMPLICATIONS

There are several implications of this research--most specifically for personnel training and selection and for human engineering. Each of these areas is related to one or more of the others, and in working out their relationships it's not easy to know just where to begin. Although the functions of equipment define the human-engineering problem, just what equipment should do is difficult to specify without the understanding of system operation that comes from intensive examination such as these RAND air-defense studies provide. And so it is with setting personnel-selection standards. They can be set once the system is analysed to see how much of



the work the equipment can take over, but here too a thorough study of system operation is needed.

However, some human-engineering and personnel requirements can be described quite generally. Once the importance of group learning is recognized it follows that equipment and facilities should be arranged not just to facilitate operation but also to help the men who operate the system learn to use its full potential most rapidly. Or, more practically, since specifying what these men are to learn is difficult unless the system can be operated under the emergency conditions it was designed for, doing anything that might hinder group learning should be avoided. Communication between members should be made as free and easy as possible. Facilities should be arranged so that each member of the group is given as complete a picture as possible of the task and how the organization is dealing with it—in central displays of some sort if these are feasible. Members of the group should be given a chance to modify their procedures. For example, the members of an air-defense crew develop priority systems for simplifying their task; if information handling is taken over by electromechanical devices, the men who run the system should be free to modify the procedures for using these devices to utilize them most efficiently.

Considering a system as an integral unit rather than as a collection of individuals says something about personnel selection. It suggests that, in manning a system, teams rather than individuals should be selected—that matching the individual to the job may be a part of the organizational development process.

The need for system training has shown up in the difficulties of getting today's complex man-machine systems to perform as expected. Reliance on

the adaptive capabilities of the operators is implicit in the design of most of these systems, but unfortunately developing these capabilities so that the system will perform adequately in an emergency requires experience under critical stresses equivalent to those of an emergency. These experiments have shown that system training, which can impose such stresses, does result in much more effective use of a system's resources—that it is one way of making the full potential of a system available before an emergency occurs. They have also enabled us to understand enough about how the organizations developed in the laboratory to formulate a useful principle that says: Train the team as a whole in an adequately simulated environment and give it knowledge of results. This technique treats the organization as a unit. It helps the organization develop by providing appropriate stress and the needed reinforcement. Although an organization gets some idea of how well it is doing just by doing it, the more complete the information about the results of its operation it gets, the more it will be reinforced. A training program, therefore, should facilitate learning by providing a factual critique which helps the organization identify its difficulties. This training principle is presently being put to use in a particular training program—the System Training Program RAND is installing in the Air Defense Command.

\* \* \*

The Systems Research Laboratory has been looking at organizations in a somewhat different way. It has considered them as integral units and, because of its choice of an organization to study, it has been able to see

the development process in some detail. Its approach seems to have quite a bit of promise for understanding and improving man-machine systems; it has many implications, not yet too specific, for integrating training, human engineering, and personnel-selection programs.

The results of the Laboratory's experiments with air-defense organizations suggest that one of the most important things in making these fantastically ingenious man-machine systems work is the use of human group-learning ability to get the greatest possible utilization of a system's resources. They also suggest that the best way to find out how to use this ability is in studies of complete systems.

The Laboratory's efforts to build a model of organizational behavior--a model that is still particular to a limited range of human behavior--may assist research workers by providing a cogent set of theoretical propositions about human behavior that can guide them in studying broader areas of human endeavor.